Adaptive h-refinement for reduced-order models

Kevin Carlberg

Sandia National Laboratories

Kevin Carlberg

Harry S. Truman Fellow, Quantitative Modeling & Analysis Department
Sandia National Laboratories, P.O. Box 969, MS 9159, Livermore, CA 94551, USA.

arXiv:1404.0442v3 [cs.NA] 18 Jul 2014

Abstract

This work presents a method to adaptively refine reduced-order models *a posteriori* without requiring additional full-order-model solves. The technique is analogous to mesh-adaptive h-refinement: it enriches the reduced-basis space online by 'splitting' a given basis vector into several vectors with disjoint support. The splitting scheme is defined by a tree structure constructed offline via recursive k-means clustering of the state variables using snapshot data. The method identifies the vectors to split online using a dual-weighted-residual approach that aims to reduce error in an output quantity of interest. The resulting method generates a hierarchy of subspaces online without requiring large-scale operations or full-order-model solves. Further, it enables the reduced-order model to satisfy *any prescribed error tolerance* regardless of its original fidelity, as a completely refined reduced-order model is mathematically equivalent to the original full-order model. Experiments on a parameterized inviscid Burgers equation highlight the ability of the method to capture phenomena (e.g., moving shocks) not contained in the span of the original reduced basis.

Keywords: adaptive refinement, h-refinement, model reduction, dual-weighted residual, adjoint error estimation, clustering

1. Introduction

Modeling and simulation of parameterized systems has become an essential tool across a wide range of industries. However, the computational cost of executing high-fidelity large-scale simulations is infeasibly high for many time-critical applications. In particular, many-query scenarios (e.g., sampling for solving statistical inverse problems) can require thousands of simulations corresponding to different input-parameter instances of the system; real-time contexts (e.g., model predictive control) require simulations to execute in mere seconds.

Reduced-order models (ROMs) have been developed to mitigate this computational bottleneck. First, they execute an 'offline' stage during which computationally expensive training tasks (e.g., evaluating the high-fidelity model at several points in the input-parameter space) compute a representative low-dimensional reduced basis for the system state. Then, during the inexpensive 'online' stage, these methods quickly compute approximate solutions for arbitrary points in the input space via a projection process of the high-fidelity full-order-model (FOM) equations onto the low-dimensional subspace spanned by the reduced basis. They

*7011 East Ave, MS 9159, Livermore, CA 94550. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy under contract DE-AC04-94-AL85000.

Email address: ktcarlb@sandia.gov (Kevin Carlberg)

URL: sandia.gov/~ktcarlb (Kevin Carlberg)

also introduce other approximations in the presence of general (i.e., not low-order polynomial) nonlinearities. See Ref. [1] and references within for a survey of current methods.

While reduced-order models almost always generate *fast* online predictions, there is no guarantee that they will generate sufficiently *accurate* online predictions. In fact, the accuracy of online predictions is predicated on the relevance of the training data to the online problem: if a physical phenomenon was not observed during the offline stage, then this feature will be missing from online predictions. In general, the most one can guarantee *a priori* is that the ROM solution error is bounded by a prescribed scalar over a finite set of 'training points' in the input-parameter space [2]. While reduced-order models can be accurate at online points contained within a reasonable neighborhood of these training points (see, e.g., Ref. [3]), they are generally inaccurate for points far outside this set.

This lack of error control[1] precludes ROMs from being employed in many contexts. For example, PDE-constrained optimization requires the solution to satisfy a prescribed forcing sequence to guarantee convergence [4]. In uncertainty quantification, if the epistemic uncertainty due to the ROM solution error dominates other sources of uncertainty, the ROM cannot be exploited in a useful manner. When simulating parameterized highly nonlinear dynamical systems, it is unlikely that any amount of training will fully encapsulate the range of complex phenomena that can be encountered online; such problems require an efficient refinement mechanism to generate accurate ROM predictions.

A few methods exist to improve a ROM solution when it is detected to be inaccurate; however, they entail *large-scale* operation counts. The most common approach is to revert to the high-fidelity model, solve the associated high-dimensional equations for the current time step or optimization iteration, add the solution to the reduced basis, and proceed with the enriched reduced-order model [5, 6, 7]. Another approach adaptively improves the reduced-order model *a posteriori* by generating a Krylov subspace [8]; here, the reduced-order model serves to accelerate the full-order solve to any specified tolerance. As our goal is to improve the reduced-order model efficiently, i.e., without incurring large-scale operations, none of these methods is appropriate.

Instead, this work proposes a novel approach inspired by mesh-adaptive h-refinement. The main idea is to adaptively refine an inaccurate ROM *online* by 'splitting' selected reduced basis vectors into multiple vectors with disjoint discrete support. This splitting technique is defined by a tree structure generated offline by applying k-means clustering to the state variables. The method uses a dual-weighted residual approach to select vectors to split online. The resulting method generates a hierarchy of subspaces online without requiring any large-scale operations or high-fidelity solves. Most importantly, the methodology acts as a 'failsafe' mechanism for the ROM: h-adaptivity enables the ROM to satisfy *any prescribed error tolerance* online, as a fully refined ROM is mathematically equivalent to the original full-order model under modest conditions.

As a final note, some 'adaptive' methods exist to tailor the ROM to specific regions of the input space [9, 10, 11, 12, 13, 14], time domain [13, 15], and state space [16, 14]. However, these methods are primarily *a priori* adaptive: they construct separate ROMs for each region *offline* with the goal of reducing the ROM dimension. While they can be used to improve the ROM *a posteriori*, e.g., by restarting the greedy algorithm online, doing so incurs additional full-order-model solves, which is what we aim to avoid.

In the remainder of this paper, matrices are denoted by capitalized bold letters, vectors by lowercase bold letters, scalars by lowercase letters, and sets by capitalized letters. The columns of a matrix $\boldsymbol{A} \in \mathbb{R}^{m \times k}$ are denoted by $\boldsymbol{a}_i \in \mathbb{R}^m$, $i \in \mathbb{N}(k)$ with $\mathbb{N}(a) := \{1, \dots, a\}$ such that $\boldsymbol{A} := [\boldsymbol{a}_1 \ \cdots \ \boldsymbol{a}_k]$. The scalar-valued matrix elements are denoted by $a_{ij} \in \mathbb{R}$ such that $\boldsymbol{a}_j := [a_{1j} \ \cdots \ a_{mj}]^T$, $j \in \mathbb{N}(k)$.

[1] Note that reduced-order-model error bounds—which exist for many problems—serve to quantify the error, while error control implies reducing this error *a posteriori*.

2. Problem formulation

2.1. Full-order model

Consider solving a parameterized sequence of systems of equations

$$\tilde{\boldsymbol{r}}^k(\boldsymbol{x}^k;\boldsymbol{\mu}) = 0 \tag{1}$$

for $k \in \mathbb{N}(t)$, where $\boldsymbol{x}^k \in \mathbb{R}^n$ denotes the state at iteration k, $\boldsymbol{\mu} \in \mathcal{D} \subset \mathbb{R}^{n_\mu}$ denotes the input parameters (e.g., boundary conditions), $\tilde{\boldsymbol{r}}^k : \mathbb{R}^n \times \mathbb{R}^{n_\mu} \to \mathbb{R}^n$ denotes the residual operator at iteration k, and t denotes maximum number of iterations. This formulation is quite general, as it describes, e.g., parameterized systems of linear equations ($t = 1$, $\tilde{\boldsymbol{r}} : (\boldsymbol{x};\boldsymbol{\mu}) \mapsto \boldsymbol{b}(\boldsymbol{\mu}) - \boldsymbol{A}(\boldsymbol{\mu})\boldsymbol{x}$) such as those arising from the finite-element discretization of elliptic PDEs, and parameterized ODEs $\dot{\boldsymbol{x}} = \boldsymbol{f}(\boldsymbol{x};\boldsymbol{\mu})$ after time discretization by an implicit linear multistep method (e.g., $\tilde{\boldsymbol{r}}^k : (\boldsymbol{x}^k;\boldsymbol{\mu}) \mapsto \boldsymbol{x}^k - \boldsymbol{x}^{k-1} - \Delta t \boldsymbol{f}\left(\boldsymbol{x}^k;\boldsymbol{\mu}\right)$ for the backward Euler scheme) such as those arising from the space- and time-discretization of parabolic and hyperbolic PDEs. Assume that we are primarily interested in computing outputs

$$z^k = g(\boldsymbol{x}^k;\boldsymbol{\mu}) \tag{2}$$

with $z^k \in \mathbb{R}$ and $g : \mathbb{R}^n \times \mathbb{R}^{n_\mu} \to \mathbb{R}$.

When the dimension n is 'large', computing the outputs of interest z^k by first solving Eq. (1) and subsequently computing outputs via Eq. (2) can be prohibitively expensive. This is particularly true for many-query (e.g., statistical inversion) and real-time (e.g., model-predictive control) problems that demand a fast evaluation of the input–output map $\boldsymbol{\mu} \mapsto \{z^1, \ldots, z^t\}$.

2.2. Reduced-order model

Model-reduction techniques aim to reduce the burden of solving Eq. (1) by employing a projection process. First, they execute a computationally expensive offline stage (e.g., solving Eq. (1) for a training set $\boldsymbol{\mu} \in \mathcal{D}_{\text{train}} \subset \mathcal{D}$) to construct 1) a low-dimensional trial basis (in matrix form) $\boldsymbol{V} \in \mathbb{R}^{n\times p}$ with $p \ll n$ that (hopefully) captures the behavior of the state $\boldsymbol{x}$ throughout the parameter domain $\mathcal{D}$, and 2) an associated test basis $\boldsymbol{W} \in \mathbb{R}^{n\times p}$. Then, during the computationally inexpensive online stage, these methods approximately solve Eq. (2) for arbitrary $\boldsymbol{\mu} \in \mathcal{D}$ by searching for solutions in the trial subspace $\bar{\boldsymbol{x}} + \text{range}\,(\boldsymbol{V}) \subset \mathbb{R}^n$ (with $\bar{\boldsymbol{x}} \in \mathbb{R}^n$ a chosen reference configuration) and enforcing the residual $\tilde{\boldsymbol{r}}^k$ to be orthogonal to the test subspace $\text{range}\,(\boldsymbol{W}) \subset \mathbb{R}^n$:

$$\boldsymbol{W}^T\tilde{\boldsymbol{r}}^k(\bar{\boldsymbol{x}} + \boldsymbol{V}\hat{\boldsymbol{x}}^k;\boldsymbol{\mu}) = 0. \tag{3}$$

Here, $\hat{\boldsymbol{x}}^k \in \mathbb{R}^p$ denotes the generalized coordinates of the reduced-order-model solution $\bar{\boldsymbol{x}}+\boldsymbol{V}\hat{\boldsymbol{x}}^k$ at iteration k. When the residual operator exhibits general nonlinear dependence on the state or is non-affine in the inputs, additional complexity-reduction approximations such as empirical interpolation [17], collocation [18, 19, 7], discrete empirical interpolation [20, 21], or gappy proper orthogonal decomposition (POD) [19, 22] are required to ensure that computing the low-dimensional residual $\boldsymbol{W}^T\tilde{\boldsymbol{r}}^k$ incurs an n-independent operation count. For simplicity, we do not consider such approximations in the present work; future work will entail extending the proposed method to such 'hyper-reduced' order models.

In many cases, the test basis can be expressed as $\boldsymbol{W} = \boldsymbol{A^n}\left(\boldsymbol{x};\boldsymbol{\mu}\right)\boldsymbol{V}$. For example, $\boldsymbol{A^n}\left(\boldsymbol{x};\boldsymbol{\mu}\right) = I$ for Galerkin projection; balanced truncation uses $\boldsymbol{A^n}\left(\boldsymbol{x};\boldsymbol{\mu}\right) = \boldsymbol{Q}$, where $\boldsymbol{Q}$ is the observability Gramian of the linear time-invariant system; the least-squares Petrov–Galerkin projection [18, 22] underlying the GNAT method employs $\boldsymbol{A^n}\left(\boldsymbol{x};\boldsymbol{\mu}\right) = \partial\tilde{\boldsymbol{r}}^k/\partial\boldsymbol{x}\left(\boldsymbol{x},\boldsymbol{\mu}\right)$; for linearized compressible-flow problems, $\boldsymbol{A^n}\left(\boldsymbol{x};\boldsymbol{\mu}\right)$ can be chosen to guarantee stability [23]. When this holds, the Petrov–Galerkin projection (3) is equivalent to a Galerkin projection performed on the modified residual $\boldsymbol{r}^k := \boldsymbol{A^n}\left(\boldsymbol{x};\boldsymbol{\mu}\right)^T\tilde{\boldsymbol{r}}^k$:

$$\boldsymbol{V}^T\boldsymbol{r}^k(\bar{\boldsymbol{x}} + \boldsymbol{V}\hat{\boldsymbol{x}}^k;\boldsymbol{\mu}) = 0, \tag{4}$$

for $k \in \mathbb{N}(p)$. In the remainder of this paper, Eq. (4) will be considered the governing equations for the reduced-order model.

2.3. Objective: adaptive refinement

The goal of this work is as follows: given a reduced basis $\boldsymbol{V}$ and online ROM solution $\hat{\boldsymbol{x}}^k$ to Eq. (4) for iteration k, 1) detect if the solution is sufficiently accurate, 2) if it is not sufficiently accurate, efficiently generate a higher-dimensional reduced basis $\boldsymbol{V}'$ with range $(\boldsymbol{V}) \subseteq$ range $(\boldsymbol{V}')$ in a goal-oriented manner that aims to reduce errors in the output z^k, 3) compute an associated solution $\hat{\boldsymbol{x}}'^k$, 4) repeat until desired accuracy is reached.

To generate this hierarchy of subspaces efficiently, we propose an analogue to adaptive h-refinement, wherein selected basis vectors $\boldsymbol{v}_i$ are 'split' online into multiple vectors with disjoint support (i.e., the element set with nonzero entries). Like all h-refinement techniques, the proposed method consists of the following components:

1. *Refinement mechanism.* In typical h-refinement, this is defined by the mesh-refinement method applied to finite elements or volumes. The proposed method refines the solution space by splitting the support of the basis vectors using a tree structure constructed via k-means clustering of the state variables. Section 3 describes this component.
2. *Error indicators.* Goal-oriented methods for h-refinement often 1) solve a coarse dual problem, 2) prolongate the adjoint solution to a representation on the fine grid, and 3) compute error estimates of the output using first-order analysis. The proposed method employs an analogous goal-oriented dual-weighted residual approach. Section 4 presents this.
3. *An adaptive algorithm.* The proposed algorithm identifies when refinement is required online and employs error indicators decide on the particular refinement, i.e., which basis vectors should be refined, and how they should be refined. Section 5 provides this algorithm.

3. Refinement mechanism

The method assumes that an initial reduced basis $\boldsymbol{V}^{(\mathbf{0})} \in \mathbb{R}^{n \times p^{(0)}}$ is provided, which is subsequently 'split' to add fidelity to the ROM online. Section 3.1 describes the tree data structure that constitutes the splitting mechanism, Section 3.2 describes how this mechanism leads to an algebraic refinement strategy, Section 3.3 highlight critical properties of the refinement method, and Section 3.4 describes construction of the tree via k-means clustering.

3.1. Tree data structure

To begin, we define a tree data structure that characterizes the refinement mechanism. The tree is characterized by a *child* function $C : \mathbb{N}(m) \to \mathcal{P}(\mathbb{N}(m))$ that describes the topology of the tree and an *element* function $E : \mathbb{N}(m) \to \mathcal{P}(\mathbb{N}(n))$ that describes the set of nonzero vector entries associated with each tree node. Here, m denotes the number of nodes in the tree and $\mathcal{P}$ denotes the powerset.

Each basis vector $\boldsymbol{v}_i$, $i \in \mathbb{N}(p)$ is characterized by a particular node on the tree $d_i \in \mathbb{N}(m)$, a set of nonzero entries (i.e., support) $E(d_i)$, and possible splits $C(d_i)$. If a given vector $\boldsymbol{v}_i$ is split, then it is replaced in the basis by $q_i := \text{card}(C(d_i))$ child vectors whose set of nonzero entries is defined by $E(k)$, $k \in C(d_i)$; the values of these nonzero entries are the same as those of the original vector $\boldsymbol{v}_i$.

We enforce the following conditions for the tree:

1. The root node includes all elements: $E(1) = \mathbb{N}(n)$, which is consistent with the possibly global support of the original reduced basis $\boldsymbol{V}^{(\mathbf{0})}$.
2. The children have disjoint support, and the union of their support equals that of the parent: For all $i \in \mathbb{N}(m)$,

$$E(j) \cap E(k) = \emptyset, \quad \forall j, k \in C(i),\ j \neq k \tag{5}$$

$$\bigcup_{j \in C(i)} E(j) = E(i). \tag{6}$$

3. Each element is associated with a single leaf node:

$$\forall l \in \mathbb{N}(n),\ \exists i \in \mathbb{N}(m) \mid E(i) = l,\ C(i) = \emptyset. \tag{7}$$

As will be shown, these requirements guarantee several critical properties of the method.

Example. Consider an example with $n = 6$ and an initial reduced basis $\boldsymbol{V}^{(\mathbf{0})} = \boldsymbol{v}_1^{(\mathbf{0})}$ of dimension 1. Figure 1 depicts an example of a tree structure for this case.

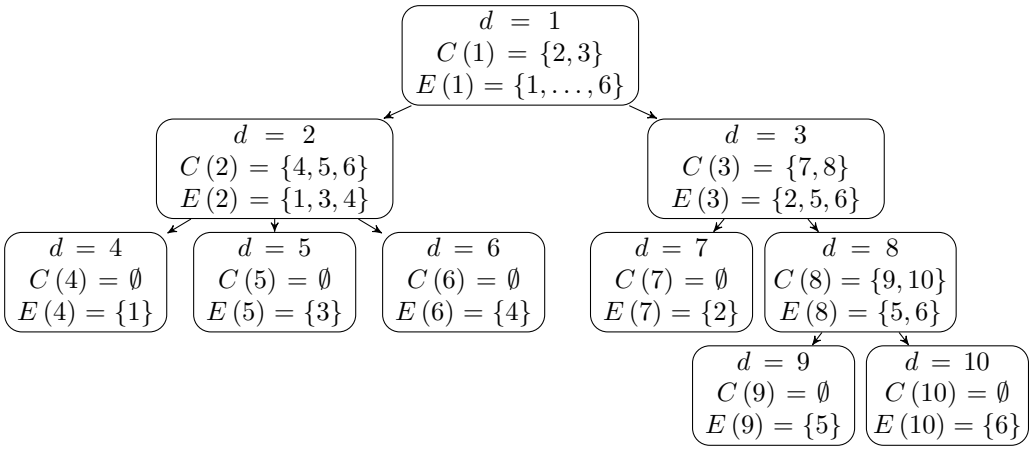

Figure 1: Tree example with $n = 6$

Suppose the basis has been split into $p = 4$ according to the tree in Figure 1 with $d_1 = 2$, $d_2 = 7$, $d_3 = 9$, and $d_4 = 10$; then, the refined reduced basis is

$$\boldsymbol{V} = \begin{bmatrix} v_{11}^{(0)} & 0 & 0 & 0 \\ 0 & v_{21}^{(0)} & 0 & 0 \\ v_{31}^{(0)} & 0 & 0 & 0 \\ v_{41}^{(0)} & 0 & 0 & 0 \\ 0 & 0 & v_{51}^{(0)} & 0 \\ 0 & 0 & 0 & v_{61}^{(0)} \end{bmatrix}. \tag{8}$$

■

In the sequel, we overload the child function for the two-argument case such that $C(i, j)$ denotes the jth child node of parent node i, where ordering of the children is implied by the binary relation $\leq$ on the natural numbers. Similarly, the overloaded element function $E(i, j)$ is the jth element for node i; again, ordering of the elements is implied by the relation $\leq$ on the natural numbers.

3.2. Refinement via basis splitting

We now put the basis-splitting methodology in the framework of typical h-refinement techniques. First, define a 'coarse' basis $\boldsymbol{V}^H \in \mathbb{R}^{n \times p}$, which is initially equal to the nominal basis $\boldsymbol{V}^{(\mathbf{0})} \in \mathbb{R}^{n \times p^{(0)}}$ with $p^{(0)} \leq p$. As this initial basis may have global support, it is characterized by $d_i = 1$, $i \in \mathbb{N}(p^{(0)})$; this is permissible due to Condition 1 of Section 3.1. Also define a 'fine' basis corresponding to the coarse basis with all vectors split according to the children of the current node. We can express the relationship between the coarse and fine bases as

$$\boldsymbol{V}^H = \boldsymbol{V}^h \boldsymbol{I}_H^h, \tag{9}$$

where $\boldsymbol{V}^h \in \mathbb{R}^{n \times q}$ with $q \geq p$ denotes the fine basis and $\boldsymbol{I}_H^h \in \{0, 1\}^{q \times p}$ denotes the *prolongation* operator. Then, for any generalized coordinates $\hat{\boldsymbol{w}}^H \in \mathbb{R}^p$ associated with the coarse basis $\boldsymbol{V}^H$, we can compute the corresponding fine representation $\hat{\boldsymbol{w}}^h \in \mathbb{R}^q$ associated with the fine basis $\boldsymbol{V}^h$ as

$$\hat{\boldsymbol{w}}_H^h = \boldsymbol{I}_H^h \hat{\boldsymbol{w}}^H, \tag{10}$$

which ensures that $\boldsymbol{V}^H\hat{\boldsymbol{w}}^H = \boldsymbol{V}^h\hat{\boldsymbol{w}}^h_H$. Note this prolongation operator is exact, unlike typical mesh-refinement strategies, where this operator is often defined as a linear or quadratic interpolant of the coarse solution on the fine grid. The *restriction* operator is not uniquely defined, but can be set, e.g., to

$$\boldsymbol{I}^H_h = \left(\boldsymbol{I}^h_H\right)^+, \tag{11}$$

where the superscript + denotes the Moore–Penrose pseudoinverse.

Using the tree structure defined in Section 3.1, we can precisely define these quantities. We first introduce the mapping $f : (i, j) \mapsto k$, which provides the fine basis-vector index k corresponding to the jth child of the ith coarse basis vector. We define it as

$$f(i, j) = \sum_{k<i} q_k + j, \quad j \in \mathbb{N}(q_i), \quad i \in \mathbb{N}(p). \tag{12}$$

In particular, note that if node d_i is a leaf (i.e., $C(d_i) = \emptyset$), then $f(i, j)$ does not exist for any j. Similarly, the inverse mapping $f^{-1} : k \mapsto (i, j)$ yields the coarse basis-vector index i and child index j corresponding to fine basis vector k.

Now, the number of vectors in the fine reduced basis is simply

$$q = \sum_{i=1}^{p} q_i. \tag{13}$$

From Condition 2 of Section 3.1, we can write the fine reduced basis as

$$v^h_{ij} = \begin{cases} v^H_{il}, & \exists k \mid j = f(l, k),\ i \in E(C(d_l, k)) \\ 0, & \text{otherwise.} \end{cases} \tag{14}$$

and the prolongation operator induced by the proposed splitting scheme as

$$[\boldsymbol{I}^h_H]_{ij} = \begin{cases} 1, & \exists k \mid i = f(j, k) \\ 0, & \text{otherwise.} \end{cases} \tag{15}$$

3.3. Properties

This section highlights several key properties of this refinement method.

Lemma 1 (Hierarchical subspaces). *The method generates a hierarchy of subspaces such that* $\text{range}\left(\boldsymbol{V}^H\right) \subseteq \text{range}\left(\boldsymbol{V}^h\right)$.

Proof *This result is self-evident from Eq.* (9)*, as*

$$\text{range}\left(\boldsymbol{V}^H\right) = \{\boldsymbol{V}^h\boldsymbol{w} \mid \boldsymbol{w} \in \text{range}\left(\boldsymbol{I}^h_H\right) \subseteq \mathbb{R}^q\} \subseteq \{\boldsymbol{V}^h\boldsymbol{w} \mid \boldsymbol{w} \in \mathbb{R}^q\} = \text{range}\left(\boldsymbol{V}^h\right). \tag{16}$$

Theorem 1 (Monotonic convergence). *If the reduced-order model* (4) *is a priori convergent, i.e., its solution satisfies*

$$\boldsymbol{V}\hat{\boldsymbol{x}}^k = \arg\min_{\boldsymbol{w} \in \text{range}(\boldsymbol{V})} \|\boldsymbol{x}^k - \bar{\boldsymbol{x}} - \boldsymbol{w}\|_{\boldsymbol{\Theta}}, \tag{17}$$

for some norm $\|\cdot\|_{\boldsymbol{\Theta}}$*, then the proposed refinement method guarantees monotonic convergence of the reduced-order-model solution, i.e.,*

$$\|\boldsymbol{x}^k - \bar{\boldsymbol{x}} - \boldsymbol{V}^h(\hat{\boldsymbol{x}}^h)^k\|_{\boldsymbol{\Theta}} \leq \|\boldsymbol{x}^k - \bar{\boldsymbol{x}} - \boldsymbol{V}^H(\hat{\boldsymbol{x}}^H)^k\|_{\boldsymbol{\Theta}}. \tag{18}$$

Proof *This follows directly from Lemma 1, as the coarse-basis solution is contained in the span of the fine basis* $\boldsymbol{V}^H(\hat{\boldsymbol{x}}^H)^k \in \text{range}\left(\boldsymbol{V}^H\right) \subseteq \text{range}\left(\boldsymbol{V}^h\right)$.

One example of a reduced-order model that satisfies the conditions of Theorem 1 arises when the residual is linear in the state and its Jacobian $\partial \boldsymbol{r}^k/\partial \boldsymbol{x}\left(\boldsymbol{\mu}\right)$ is symmetric and positive definite. In this case, a Galerkin-projection ROM satisfies Eq. (17) for $\boldsymbol{\Theta} = \partial \boldsymbol{r}^k/\partial \boldsymbol{x}\left(\boldsymbol{\mu}\right)$ with $\|\boldsymbol{w}\|_{\partial \boldsymbol{r}^k/\partial \boldsymbol{x}(\boldsymbol{\mu})} := \sqrt{\boldsymbol{w}^T \partial \boldsymbol{r}^k/\partial \boldsymbol{x}\left(\boldsymbol{\mu}\right) \boldsymbol{w}}$. Another example is least-squares Petrov–Galerkin applied to a parametrized system of linear equations [22], where $\boldsymbol{\Theta} = \left(\partial \boldsymbol{r}^k/\partial \boldsymbol{x}\left(\boldsymbol{\mu}\right)\right)^T \partial \boldsymbol{r}^k/\partial \boldsymbol{x}\left(\boldsymbol{\mu}\right)$.

Theorem 2 (Convergence to the full-order model). *If every element has a nonzero entry in one of the original reduced-basis vectors, i.e.,*

$$\forall l \in \mathbb{N}(n),\ \exists\ (i,j) \in \mathbb{N}(n) \times \mathbb{N}(p^{(0)}) \mid v_{ij}^{(0)} \neq 0, \tag{19}$$

and Eq. (7) *holds, then a completely split basis yields a reduced-order model equivalent to the full-order model.*

Proof *Under these conditions, a completely split basis can be written as* $\boldsymbol{V} \in \mathbb{R}^{n \times np^{(0)}}$ *with all basis vectors in the leaf-node state, i.e.,* $C\left(d_i\right) = \emptyset$, $i \in \mathbb{N}(np^{(0)})$. *Because Eq.* (7) *guarantees that each element is associated with a single leaf node, this implies that*

$$\forall l \in \mathbb{N}(n),\ \exists\ i \in \mathbb{N}(np^{(0)}) \mid \boldsymbol{v}_i = \boldsymbol{e}_l \beta_i, \tag{20}$$

where $\boldsymbol{e}_l \in \{0,1\}^n$ *denotes the* l*th canonical unit vector and* $\beta_i \neq 0$, $i \in \mathbb{N}(np^{(0)})$. *Eq.* (20) *implies that the completely split basis can be post-multiplied by a (weighted) permutation matrix to yield the* $n \times n$ *identity matrix* $\boldsymbol{I}_n$, *i.e.,*

$$\boldsymbol{I}_n = \boldsymbol{V}\boldsymbol{\Gamma}. \tag{21}$$

Here, the matrix $\boldsymbol{\Gamma} \in \mathbb{R}^{np^{(0)} \times n}$ *consists of columns*

$$\boldsymbol{\gamma}_l = \frac{1}{\beta_i}\boldsymbol{e}_i, \quad i \in \{j \mid \boldsymbol{v}_j = \boldsymbol{e}_l \beta_j\}, \quad l \in \mathbb{N}(n). \tag{22}$$

Eq. (21) *implies that*

$$\text{range}\left(\boldsymbol{I}_n\right) = \mathbb{R}^n \subseteq \text{range}\left(\boldsymbol{V}\right) \subseteq \mathbb{R}^n, \tag{23}$$

which completes the proof.

Lemma 1 and Theorem 2 show that the proposed refinement method enables the reduced-order model to generate a sequence of hierarchical subspaces that converges to the full-order model under modest assumptions. Thus, the method acts as a 'failsafe' mechanism: it allows the reduced-order model to generate *arbitrarily accurate solutions*. Despite this result, the associated rate of convergence is unknown, which precludes any *a priori* guarantee that the h-adaptive ROM will remain truly low dimensional for stringent accuracy requirements. However, numerical experiments in Section 6 demonstrate that the proposed method often leads to accurate responses with low-dimensional refined bases.

Remark. Note that the refinement method does not preclude a rank-deficient basis; this can be seen from Theorem 2, wherein a completely split basis has $np^{(0)} \geq n$ columns. To detect (and remove) rank deficiency, the refinement algorithm computes a rank-revealing QR factorization after each split (Steps 14–15 of Algorithm 4 and Steps 28–29 of Algorithm 5). ■

3.4. Tree construction via k-means clustering of the state variables

Any tree that satisfies Conditions 1–3 of Section 3.1 will lead to the critical properties proved in Section 3.3. This section presents one such tree-construction approach, which executes offline and employs the following heuristic:

> State variables x_i that tend to be strongly positively or negatively correlated can be accurately represented by the same generalized coordinate, and should therefore reside in the same tree node.

Example. To justify this heuristic, consider an example with $n = 6$ degrees of freedom and $n_o = 8$ observations of the state, e.g., from a computed time history. Assume that snapshots can be decomposed as

$$\boldsymbol{X} = \sum_{i=1}^{3} \boldsymbol{y}_i \boldsymbol{z}_i^T + 0.1\boldsymbol{E} \tag{24}$$

where $\boldsymbol{E} \in [-1, 1]^{n \times n_o}$ is a matrix of random uniformly distributed noise and the data matrices are

$$\boldsymbol{Z} = \begin{bmatrix} -2.2083 & -5.1072 & 2.6816 & 9.3277 & -6.4506 & -3.2548 & 4.2237 & -3.2557 \\ -2.9810 & 0.6557 & 3.0474 & 5.5252 & 2.7674 & 2.3311 & 9.6190 & -6.6484 \\ -2.4547 & 5.2676 & -3.6434 & 5.5661 & -7.5449 & 9.3079 & -2.0459 & -0.0728 \end{bmatrix}^T$$

$$\boldsymbol{Y} = \begin{bmatrix} -3.9885 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 8.6843 & 0 & 0 & -1.6393 \\ 0 & -1.7288 & 0 & 6.0559 & 2.2407 & 0 \end{bmatrix}^T.$$

The sparsity structure of $\boldsymbol{Y}$ implies that the following sets of state variables are strongly correlated or anti-correlated across observations: $\{1\}$, $\{3, 6\}$, and $\{2, 4, 5\}$. This is apparent from computing the matrix of sample correlation coefficients:

$$\boldsymbol{R} = \begin{bmatrix} 1.0000 & 0.1526 & -0.5698 & -0.1534 & -0.1554 & 0.5705 \\ 0.1526 & 1.0000 & -0.0180 & -1.0000 & -1.0000 & 0.0198 \\ -0.5698 & -0.0180 & 1.0000 & 0.0209 & 0.0212 & -1.0000 \\ -0.1534 & -1.0000 & 0.0209 & 1.0000 & 1.0000 & -0.0227 \\ -0.1554 & -1.0000 & 0.0212 & 1.0000 & 1.0000 & -0.0229 \\ 0.5705 & 0.0198 & -1.0000 & -0.0227 & -0.0229 & 1.0000 \end{bmatrix}. \tag{25}$$

Suppose we start with a one-dimensional reduced basis corresponding to the first left singular vector of $\boldsymbol{X}$

$$\boldsymbol{V}^{(0)} = \boldsymbol{V}^H = \boldsymbol{v}_1^H = \begin{bmatrix} -0.2609 & -0.0348 & 0.9390 & 0.1240 & 0.0463 & -0.1773 \end{bmatrix}^T.$$

Because the data nearly lie in a three-dimensional subspace of $\mathbb{R}^6$, the optimal performance of a refinement scheme would yield small error after splitting this one-dimensional basis into a basis of dimension three. Thus, consider splitting $\boldsymbol{V}^H$ into three children using a tree that follows the stated heuristic, i.e., is characterized by $C(1) = \{2, 3, 4\}$, $E(2) = \{1\}$, $E(3) = \{3, 6\}$, and $E(4) = \{2, 4, 5\}$. The resulting basis becomes

$$\boldsymbol{V}^h = \begin{bmatrix} -0.2609 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0.9390 & 0 & 0 & -0.1773 \\ 0 & -0.0348 & 0 & 0.1240 & 0.0463 & 0 \end{bmatrix}^T.$$

The resulting projection error of the data is merely $\|\boldsymbol{X} - \boldsymbol{V}^h \left(\boldsymbol{V}^h\right)^+ \boldsymbol{X}\|_F / \|\boldsymbol{X}\|_F = 0.0033$. By contrast, generating an alternative three-dimensional fine basis $\bar{\boldsymbol{V}}^h$ by splitting the basis using a (similar) tree characterized by $E(2) = \{1\}$, $E(3) = \{3, 5\}$, $E(3) = \{2, 4, 6\}$, yields a much larger error of $\|\boldsymbol{X} - \bar{\boldsymbol{V}}^h \left(\bar{\boldsymbol{V}}^h\right)^+ \boldsymbol{X}\|_F / \|\boldsymbol{X}\|_F = 0.4948$.

One way to identify these correlated variables is to employ k-means clustering [24] after pre-processing the data by 1) normalizing observations of each variable (to enable clustering to detect correlation), and 2) negating the observation vector if the first observation is negative (to enable clustering to detect anti-correlation). This is visualized in Figure 2 for the current example. Note that correlated and anti-correlated variables have a small Euclidean distance between them after this processing; this allows k-means clustering to identify them as a group. ■

To this end, we construct the tree offline by recursively applying k-means clustering to observations of the state variables (after reference subtraction, normalization, and origin flipping). Algorithm 1 describes the

Algorithm 1 Tree construction via recursive k-means clustering (offline)

Input: n_o snapshots of the reference-centered[2] state in matrix form $\boldsymbol{X} \in \mathbb{R}^{n \times n_o}$, number of means $\bar{k}$
Output: child function C, element function E, and number of nodes m

for $i = 1, \ldots, n$ **do**
  Normalize rows of $\boldsymbol{X}$ to capture correlation by clustering $\boldsymbol{x}_i^T \leftarrow \boldsymbol{x}_i^T / \|\boldsymbol{x}_i^T\|$
  if $x_{i1} < 0$ **then** {Flip over origin to capture negative correlation by clustering}
    $\boldsymbol{x}_i^T \leftarrow -\boldsymbol{x}_i^T$
  end if
end for
Set root node to contain all elements $E(1) = \mathbb{N}(n)$.
Initialize recent-node set $D \leftarrow \{1\}$ and node count $m \leftarrow 1$.
while $\text{card}(D) > 0$ **do**
  $\bar{D} \leftarrow D$, $D \leftarrow \emptyset$
  for $i = 1, \ldots, \text{card}\left(\bar{D}\right)$ **do**
    Set splitting node to the ith element of the recent-node set $d \leftarrow \bar{D}(i)$, where ordering is implied by $\geq$ on the natural numbers.
    if $E(d) = \emptyset$ **then** {No elements to split}
      Continue
    end if
    Select snapshots of current elements $\bar{x}_{jk} \leftarrow x_{E(d,j)k}$, $j \in \mathbb{N}(\text{card}(E(d)))$, $k \in \mathbb{N}(n_o)$
    $(\bar{E}_1, \ldots, \bar{E}_{n_c}) = \text{kmeans}\left(\bar{\boldsymbol{X}}, \bar{k}\right)$, where $\bar{E}_j \subset \mathbb{N}(\text{card}(E(d)))$ denotes the set of elements in cluster j, and n_c denotes the number of non-empty clusters.
    if $n_c = 1$ **then** {Cannot have only one child}
      for $j = 1, \ldots, \text{card}(E(d))$ **do** {Make all children into leaf nodes}
        $\bar{E}_j = j$
      end for
    end if
    for $j = 1, \ldots, n_c$ **do**
      $m \leftarrow m + 1$
      $D \leftarrow D \cup m$
      $E(m) = \{E(d, j) \mid j \in \bar{E}_j\}$
      $C(d, j) = m$
    end for
  end for
end while

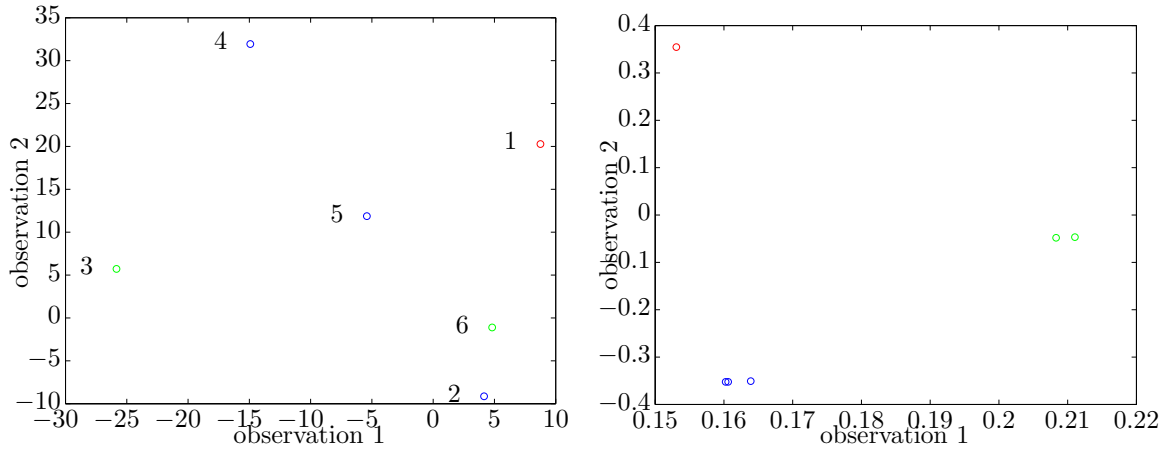

(a) before processing (state variables labeled) (b) after processing

Figure 2: First two observations of the state variables (i.e., first two columns of $\boldsymbol{X}$) for the example in Section 3.4. After processing these observations by normalization and origin flipping, correlated and anti-correlated state variables are separated by small geometric distances and can thus be grouped via clustering.

method. The n_o observations of these variables are obtained from snapshot data, which are often available, e.g., when the reduced basis is constructed via proper orthogonal decomposition.

4. Dual-weighted residual error indicators

To compute error indicators for refinement, we propose a goal-oriented dual-weighted residual methodology based on adjoint solves. It can be considered a model-reduction adaptation of duality-based error-control methods developed for differential equations [25, 26], finite-element discretizations [27, 28, 29, 30], finite-volume discretizations [31, 32, 33], and discontinuous Galerkin discretizations [34, 35]. Because the proposed method performs refinement online at the iteration level, it requires error indicators associated with the error in ROM output at iteration k, i.e., $g(\bar{\boldsymbol{x}}+\boldsymbol{V}\hat{\boldsymbol{x}}^k;\boldsymbol{\mu})$. To simplify notation in this section, we set $\bar{\boldsymbol{x}}=0$ and write the associated single solve (Eq. (4) for a single iteration and parameter instance) simply as

$$\boldsymbol{V}^T\boldsymbol{r}(\boldsymbol{V}\hat{\boldsymbol{x}})=0. \tag{26}$$

First, we approximate the output due to the (unknown) fine solution $\hat{\boldsymbol{x}}^h$ to first-order about the coarse solution $\hat{\boldsymbol{x}}^H$:

$$g\left(\boldsymbol{V}^h\hat{\boldsymbol{x}}^h\right)\approx g\left(\boldsymbol{V}^H\hat{\boldsymbol{x}}^H\right)+\frac{\partial g}{\partial \boldsymbol{x}}\left(\boldsymbol{V}^H\hat{\boldsymbol{x}}^H\right)\boldsymbol{V}^h\left(\hat{\boldsymbol{x}}^h-\boldsymbol{I}_H^h\hat{\boldsymbol{x}}^H\right), \tag{27}$$

where we have used Eq. (9) to relate the coarse and fine bases. Similarly, we can approximate the fine residual to first order about the coarse solution as

$$0=(\boldsymbol{V}^h)^T\boldsymbol{r}\left(\boldsymbol{V}^h\hat{\boldsymbol{x}}^h\right)\approx(\boldsymbol{V}^h)^T\boldsymbol{r}\left(\boldsymbol{V}^H\hat{\boldsymbol{x}}^H\right)+(\boldsymbol{V}^h)^T\frac{\partial \boldsymbol{r}}{\partial \boldsymbol{x}}\left(\boldsymbol{V}^H\hat{\boldsymbol{x}}^H\right)\boldsymbol{V}^h\left(\hat{\boldsymbol{x}}^h-\boldsymbol{I}_H^h\hat{\boldsymbol{x}}^H\right). \tag{28}$$

Solving for the state error yields

$$\left(\hat{\boldsymbol{x}}^h-\boldsymbol{I}_H^h\hat{\boldsymbol{x}}^H\right)\approx-\left[(\boldsymbol{V}^h)^T\frac{\partial \boldsymbol{r}}{\partial \boldsymbol{x}}\left(\boldsymbol{V}^H\hat{\boldsymbol{x}}^H\right)\boldsymbol{V}^h\right]^{-1}(\boldsymbol{V}^h)^T\boldsymbol{r}\left(\boldsymbol{V}^H\hat{\boldsymbol{x}}^H\right) \tag{29}$$

[2]This implies that the reference state $\bar{\boldsymbol{x}}$ should be subtracted from the state snapshots.

Substituting (29) in (27) yields

$$g\left(\boldsymbol{V}^h\hat{\boldsymbol{x}}^h\right) - g\left(\boldsymbol{V}^H\hat{\boldsymbol{x}}^H\right) \approx -\left(\hat{\boldsymbol{y}}^h\right)^T(\boldsymbol{V}^h)^T\boldsymbol{r}\left(\boldsymbol{V}^H\hat{\boldsymbol{x}}^H\right). \tag{30}$$

where the fine adjoint solution $\hat{\boldsymbol{y}}^h \in \mathbb{R}^q$ satisfies

$$(\boldsymbol{V}^h)^T\frac{\partial \boldsymbol{r}^k}{\partial \boldsymbol{x}}(\boldsymbol{V}^H\hat{\boldsymbol{x}}^H)^T\boldsymbol{V}^h\hat{\boldsymbol{y}}^h = (\boldsymbol{V}^h)^T\frac{\partial g}{\partial \boldsymbol{x}}\left(\boldsymbol{V}^H\hat{\boldsymbol{x}}^H\right)^T. \tag{31}$$

Because we would like to avoid q-dimensional solves associated with the fine basis $\boldsymbol{V}^h$, we approximate $\hat{\boldsymbol{y}}^h$ as the prolongation of the coarse adjoint solution

$$\hat{\boldsymbol{y}}^h_H = \boldsymbol{I}^h_H\hat{\boldsymbol{y}}^H, \tag{32}$$

where $\hat{\boldsymbol{y}}^H$ satisfies

$$\left(\boldsymbol{V}^H\right)^T\frac{\partial \boldsymbol{r}^k}{\partial \boldsymbol{x}}(\boldsymbol{V}^H\hat{\boldsymbol{x}}^H)^T\boldsymbol{V}^H\hat{\boldsymbol{y}}^H = \left(\boldsymbol{V}^H\right)^T\frac{\partial g}{\partial \boldsymbol{x}}\left(\boldsymbol{V}^H\hat{\boldsymbol{x}}^H\right)^T \tag{33}$$

Substituting the approximation $\hat{\boldsymbol{y}}^h_H$ for $\hat{\boldsymbol{y}}^h$ in (30) yields a cheaply computable error estimate

$$g\left(\boldsymbol{V}^h\hat{\boldsymbol{x}}^h\right) - g\left(\boldsymbol{V}^H\hat{\boldsymbol{x}}^H\right) \approx -\left(\hat{\boldsymbol{y}}^h_H\right)^T(\boldsymbol{V}^h)^T\boldsymbol{r}\left(\boldsymbol{V}^H\hat{\boldsymbol{x}}^H\right). \tag{34}$$

The right-hand side can be bounded as

$$|\left(\hat{\boldsymbol{y}}^h_H\right)^T(\boldsymbol{V}^h)^T\boldsymbol{r}\left(\boldsymbol{V}^H\hat{\boldsymbol{x}}^H\right)| \leq \sum_{i\in\mathbb{N}(q)}\delta^h_i, \tag{35}$$

where the error indicators $\delta^h_i \in \mathbb{R}_+$, $i \in \mathbb{N}(q)$ are

$$\delta^h_i = |\left[\hat{y}^h_H\right]_i\left(\boldsymbol{v}^h_i\right)^T\boldsymbol{r}\left(\boldsymbol{V}^H\hat{\boldsymbol{x}}^H\right)|. \tag{36}$$

Meyer and Matties [36] also proposed a dual-weighted residual method for reduced-order models. However, their approach was not applied to adaptive refinement and did not consider a hierarchy of reduced bases; further, their proposed dual solve was carried out on the full-order model, which is infeasibly expensive for the present context.

Remark. Some mesh-refinement techniques [31, 32] advocate computing refinement indicators that minimize the error in the computable correction

$$\left(\hat{\boldsymbol{y}}^h - \hat{\boldsymbol{y}}^h_H\right)^T(\boldsymbol{V}^h)^T\boldsymbol{r}\left(\boldsymbol{V}^H\hat{\boldsymbol{x}}^H\right).$$

To approximate this quantity, they employ prolongation operators of varying fidelity, e.g., linear and quadratic interpolants. Such a strategy is not straightforwardly applicable to the current context, as the prolongation operator $\boldsymbol{I}^h_H$ is exact. ■

5. Adaptive *h*-refinement algorithm

We now return to the original objective of this paper: adaptively refine the reduced-order model online. Algorithm 3 describes our proposed methodology for achieving this within a time-integration scheme. Step 1 first computes the reduced-order-model solution satisfying a tolerance ϵ_{ROM}. Then in Step 2, refinement occurs if the norm of the *full-order residual* is above a desired threshold ϵ. Note that other (inexpensive) error indicators could be used to flag refinement, e.g., error surrogates [37]. Refinement continues until this full-order tolerance is satisfied; note that any tolerance can be reached, as a completely split basis yields a reduced-order model equivalent to the full-order model (see Section 3.1). Finally, Step 7 resets the basis

Algorithm 2 Error estimates (online)

Input: coarse reduced basis $\boldsymbol{V}^H$, coarse solution $\hat{\boldsymbol{x}}^H$
Output: fine reduced basis $\boldsymbol{V}^h$, fine error-estimate vector $\boldsymbol{\delta}^h$
1: Solve coarse adjoint problem (33) for $\hat{\boldsymbol{y}}^H$.
2: Define prolongation operator $\boldsymbol{I}_H^h$ via Eq. (15).
3: Define fine reduced basis $\boldsymbol{V}^h$ via Eq. (9) and fine representation of adjoint solution $\hat{\boldsymbol{y}}_H^h$ via Eq. (32)
4: Compute fine error-estimate vector $\boldsymbol{\delta}^h$ via Eq. (36)

Algorithm 3 Adaptive h-refinement (online)

Input: iteration k, basis $\boldsymbol{V}$, ROM solver tolerance ϵ_{ROM}, FOM solver tolerance ϵ
Output: updated basis $\boldsymbol{V}$, generalized state $\hat{\boldsymbol{x}}^k$
1: Compute ROM solution $\hat{\boldsymbol{x}}^k$ satisfying $\|\boldsymbol{V}^T\boldsymbol{r}^k(\bar{\boldsymbol{x}} + \boldsymbol{V}\hat{\boldsymbol{x}}^k;\boldsymbol{\mu})\| \leq \epsilon_{\text{ROM}}$.
2: **if** FOM not converged $\|\boldsymbol{r}^k(\bar{\boldsymbol{x}} + \boldsymbol{V}\hat{\boldsymbol{x}}^k;\boldsymbol{\mu})\| > \epsilon$ **then**
3: Refine basis via Algorithm 4: $\boldsymbol{V} \leftarrow \text{Refine}\left(\boldsymbol{V}, \hat{\boldsymbol{x}}^k\right)$.
4: Return to Step 1.
5: **end if**
6: **if** $\mod(k, n_{\text{reset}}) = 0$ **then**
7: Reset basis $\boldsymbol{V} \leftarrow \boldsymbol{V}^{(\mathbf{0})}$.
8: **end if**

every n_{reset} time iterations. This ensures 1) the basis does not grow monotonically, and 2) work performed to refine the basis can be amortized over subsequent time steps, where the solution is unlikely to significantly change. Note that if Step 1 entails an iterative solve (e.g., Newton), then the pre-refinement solution can be employed as an initial guess.

Algorithm 4 describes the proposed method for refining the basis using the refinement mechanism and error indicators presented in Sections 3 and 4, respectively. Appendix Appendix A describes a more sophisticated approach wherein the basis vectors are not split into all possible children; the children are separated into groups, each of which contributes roughly the same fraction of that vector's error.

First, Step 1 of Algorithm 4 computes error estimates for the fine basis (i.e., current basis with all vectors split into all possible children) using the dual-weighted residual approach. Step 3 marks the parent basis vectors to refine: those with above-average error contribution from its children. Steps 5–8 split the parent vector i into vectors corresponding to its q_i children according to the defined tree. Steps 9–12 update the reduced basis and tree nodes. Because this split does not guarantee a full-ranks basis, Step 14 performs an efficient QR factorization with column pivoting to identify 'redundant' basis vectors. Step 15 subsequently removes these vectors from the basis and Step 16 performs the necessary bookkeeping for the tree nodes.

6. Numerical experiments: parameterized inviscid Burgers' equation

We assess the method's performance on the parameterized inviscid Burgers' equation. While simple, this problem is particularly challenging for reduced-order models. This arises from the fact that ROMs approximate the solution as a linear combination of spatially fixed reduced-basis functions; as such, they work well when the dynamics are primarily Eulerian, i.e., are fixed with respect to the underlying grid. However, when the dynamics are Lagrangian in nature and exhibit motion with respect to the underlying grid (e.g., moving shocks), reduced-order models generally fail to capture the critical phenomenon at every time step and parameter instance.

We employ the problem setup described in Ref. [38]. Consider the parameterized initial boundary value

Algorithm 4 Refine (online)

Input: initial basis $\boldsymbol{V}$, reduced solution $\hat{\boldsymbol{x}}$
Output: refined basis $\boldsymbol{V}$

```
Compute fine error-estimate vector and fine reduced basis via Algorithm 2:
(δ^h, V^h) ← Error estimates (V, x̂).
Put local error estimates in parent–child format η_ij = δ^h_{f(i,j)}, i ∈ ℕ(p), j ∈ ℕ(q_i).
Mark basis vectors to refine I = {i | Σ_j η_ij ≥ 1/p Σ_kj η_kj}
for i ∈ I do {Split v_i into q_i vectors}
   for k ∈ ℕ(q_i) do
      x_k = v^h_{f(i,k)}
      d̄_k = C (d_i, k)
   end for
   v_i ← x_1, d_i ← d̄_1
   for k = 2, ... q_i do
      v_{p+k-1} ← x_k, d_{p+k-1} ← d̄_k,
   end for
end for
Compute thin QR factorization with column pivoting V = QR, RΠ̄ = Q̄R̄.
Ensure full-rank matrix V ← V [π̄_1 ··· π̄_r], where r denotes the numerical rank of R.
Update tree [d_1 ··· d_r] ← [d_1 ··· d_p] [π̄_1 ··· π̄_r].
```

problem

$$\frac{\partial u(x,\tau)}{\partial \tau} + \frac{1}{2}\frac{\partial \left(u^2\left(x,\tau\right)\right)}{\partial x} = 0.02e^{\mu_2 x} \tag{37}$$

$$u(0,\tau) = \mu_1,\ \forall \tau > 0 \tag{38}$$

$$u(x,0) = 1,\ \forall x \in [0,\ 100]\,, \tag{39}$$

where μ_1 and μ_2 are two real-valued input variables. Godunov's scheme discretizes the problem, which leads to a finite-volume formulation consistent with the original formulation in Eq. (1). The one-dimensional domain is discretized using a grid with 251 nodes corresponding to coordinates coordinates $x_i = i \times (100/250)$, $i = 0, \ldots, 250$. Hence, the resulting full-order model is of dimension $n = 250$. The solution $u(x,\tau)$ is computed in the time interval $\tau \in [0, 50]$ using a uniform computational time-step size $\Delta t = 0.05$, leading to $t = 1000$ total time steps.

For simplicity, we employ a POD–Galerkin ROM. During the offline stage, snapshots of the state are collected for the first t_{train} time steps at training inputs. Then, the initial condition is subtracted from these snapshots, and they are concatenated column-wise to generate the snapshot matrix. Finally, the thin singular value decomposition of the snapshot matrix is computed, and the initial reduced basis $\boldsymbol{V}^{(\mathbf{0})}$ is set to the first $p^{(0)}$ left singular vectors. During the online stage, a Galerkin projection is employed using this reduced basis. For all experiments, the initial condition is set to the reference condition, i.e., $\bar{\boldsymbol{x}} = \boldsymbol{x}^0$. For h-adaptivity, we set the number of means to $\bar{k} = 10$ in Algorithm 1. For Algorithm 2, the output of interest is set to the residual norm, i.e., $g(\boldsymbol{x}^k; \boldsymbol{\mu}) = \|\tilde{\boldsymbol{r}}^k(\boldsymbol{x}^k; \boldsymbol{\mu})\|_2^2$. For Algorithm 3, the ROM tolerance is set to $\epsilon_{\text{ROM}} = 5 \times 10^{-3}$.[3] The basis-reset frequency n_{reset} will vary during the experiments. Step 1 incurs a Newton solve; when refinement has occurred, the initial guess is set to the converged solution from the previous refinement level. Finally, the experiments employ the (more complex) Refine method defined by Algorithm 5 with a child-partition factor $\alpha = 2$.

Note that because the residual operator is nonlinear in the state, a projection alone is insufficient to generate computational savings over the full-order model. Future work will address extending the proposed h-

[3]For the ROMs without adaptivity, the ROM convergence tolerance is set to $\epsilon_{\text{ROM}} = 1 \times 10^{-5}$.

refinement method to ROMs equipped with a complexity reduction mechanism such as empirical interpolation or gappy POD.

6.1. Fixed inputs

For this example, the input parameters are set to $\mu_1 = 3$ and $\mu_2 = 0.02$. However, the problem can be considered to be predictive, as we only collect snapshots in the time interval $\tau_{\text{train}} \in [0, 7.5]$, i.e., for the first $t_{\text{train}} = 150$ time steps. This choice is made to introduce a significant challenge for the ROM: while the (unrefined) reduced basis captures discontinuities that arise in the first 150 time steps, it will not capture such discontinuities that arise outside of this time interval.[4]

Table 1 reports results for typical POD–Galerkin ROMs of differing dimensions, as well as results for the proposed h-refinement method with different parameters and a FOM tolerance in Algorithm 3 of $\epsilon = 0.05$. Here, the relative error is defined as

$$\text{relative error} = \frac{1}{t}\sum_{k=1}^{t} \|u_{\text{FOM}}(\cdot, \tau^k) - u_{\text{ROM}}(\cdot, \tau^k)\|_{L_2} / \|u_{\text{FOM}}(\cdot, \tau^k)\|_{L_2}.$$

Figure 3 compares the solutions predicted by POD–Galerkin with no basis truncation (i.e., $p = 150$) and that of the proposed method with an initial basis size of $p^{(0)} = 10$ with $\boldsymbol{V}^{(\mathbf{0})} \in \mathbb{R}^{n \times p^{(0)}}$ and a basis-reset frequency of $n_{\text{reset}} = 50$.

	no adaptivity			h-adaptivity				
initial basis dimension $p^{(0)}$	10	45	150	5	10	20	10	10
basis-reset frequency n_{reset}				50	50	50	100	25
average basis dimension per Newton iteration $\bar{p}$	10	45	150	41.4	44.3	58	73	37
average number of Refine calls per time step				0.20	0.19	0.14	0.13	0.28
relative error (%)	45.8	43.9	8.5	0.3	0.5	0.2	0.2	0.3
online time (seconds)	1.4	2.14	5.77	5.53	4.63	7.27	6.90	7.46

Table 1: Comparison between POD–Galerkin ROMs without refinement and with h-adaptive refinement for the fixed-inputs case.

First, note that the reduced-order model is highly inaccurate (even when the basis is not truncated) unless equipped with h-adaptivity. The reason for this is simple: the training has not captured the flow regime with shock locations past approximately $x = 60$. This illustrates a powerful capability of the proposed h-adaptation methodology: it enables ROMs to be incrementally refined to capture previously unobserved phenomena. In fact, the average basis dimension (per Newton iteration) for the best-performing h-adaptive ROM ($p^{(0)} = 10$, $n_{\text{reset}} = 50$) is only $\bar{p} = 44.3$, which is smaller than the basis dimensions for ROMs without adaptivity ($p = 45$ and $p = 150$) that yield much higher errors (43.9% and 8.5%, respectively).

Second, adaptation parameters $p^{(0)}$ and n_{reset} both lead to a performance tradeoff. When $p^{(0)}$ is small, it leads to smaller average basis sizes $\bar{p}$. However, it increases the number of Refine calls per time step, as the smaller basis must be refined more times to achieve desired accuracy. Similarly, resetting the basis more frequently (smaller n_{reset}) leads to a smaller $\bar{p}$, but more average refinement steps. As such, an intermediate value of both parameters leads to the shortest online evaluation time.

Finally, notice that the online evaluation time for the adaptive ROM with an average basis size of $\bar{p} = 44.3$ is roughly twice that of a non-adaptive ROM with roughly the same basis size $p = 45$. This discrepancy in evaluation time can be attributed to the overhead in performing the adaptation. For larger problem sizes, one would expect this overhead to be smaller relative to the total online evaluation time.

[4]Note that the refinement method can also be applied when the original reduced basis captures all relevant online phenomena; however, the need for *a posteriori* refinement is weaker in this case.

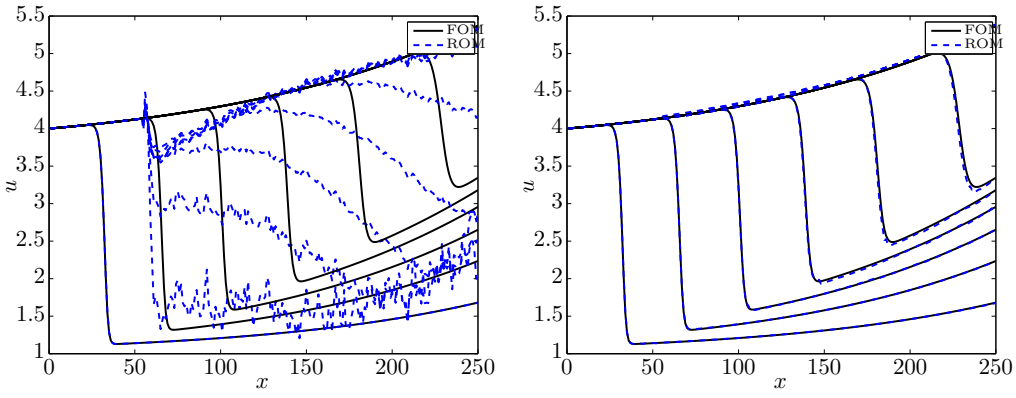

(a) no adaptivity, no truncation ($p = 150$) (b) h-adaptivity, $p^{(0)} = 10$, $n_{\text{reset}} = 50$

Figure 3: Comparison of solutions computed by POD–Galerkin with and without h-adaptivity for the fixed-inputs case.

Next, we assess the performance of the h-refinement method as the full-order-model tolerance ϵ in Algorithm 3 varies. Table 2 and Figure 4 report the results. As expected, the proposed method allows the ROM to achieve any of the prescribed tolerances. As the tolerance becomes more rigorous, the ROM solution improves; however, it does so at increased computational cost, as both the average basis dimension $\bar{p}$ and number of Refine calls per time step increase to satisfy the requirement.

	$\epsilon = 0.35$	$\epsilon = 0.05$	$\epsilon = 0.01$
average basis dimension per Newton iteration $\bar{p}$	33.6	44.2507	53.9
average number of Refine calls per time step	0.115	0.189	0.212
relative error (%)	12.2	0.51	0.078
online time (seconds)	4.61	4.63	7.64

Table 2: Effect of full-order-model tolerance ϵ on h-adaptive refinement for $p^{(0)} = 10$ and $n_{\text{reset}} = 50$ for the fixed-inputs case.

6.2. Input variation

For this experiment, we assess the proposed methodology in an input-varying scenario. In particular, the offline stage collects snapshots in the time interval $\tau_{\text{train}} \in [0, 2.5]$ for the training set $\{\boldsymbol{\mu}^1, \ldots, \boldsymbol{\mu}^3\}$ described in Table 3, which is constructed by uniformly sampling the input space along $(\mu_1, \mu_2) = (3\alpha, 0.02\alpha)$, $\alpha \in [1, 3]$.

Figure 5 and Table 4 report the results for this experiment. The same phenomena are prevalent as were apparent in the previous experiment. The primary difference is that the POD–Galerkin model without adaptivity performs better than previously (due to more informative snapshots). However, h-adaptivity is still required to drive errors below 1%. Note that the proposed method compensated for an unsophisticated uniform-sampling of the input space. The method would still be applicable for more rigorous (e.g., POD–

Table 3: Offline and online inputs for the inviscid Burgers equation

Input variables	Training point $\boldsymbol{\mu}^1$	Training point $\boldsymbol{\mu}^2$	Training point $\boldsymbol{\mu}^3$	Online point $\boldsymbol{\mu}^\star$
μ_1	3	6	9	4.5
μ_2	0.02	0.05	0.075	0.038

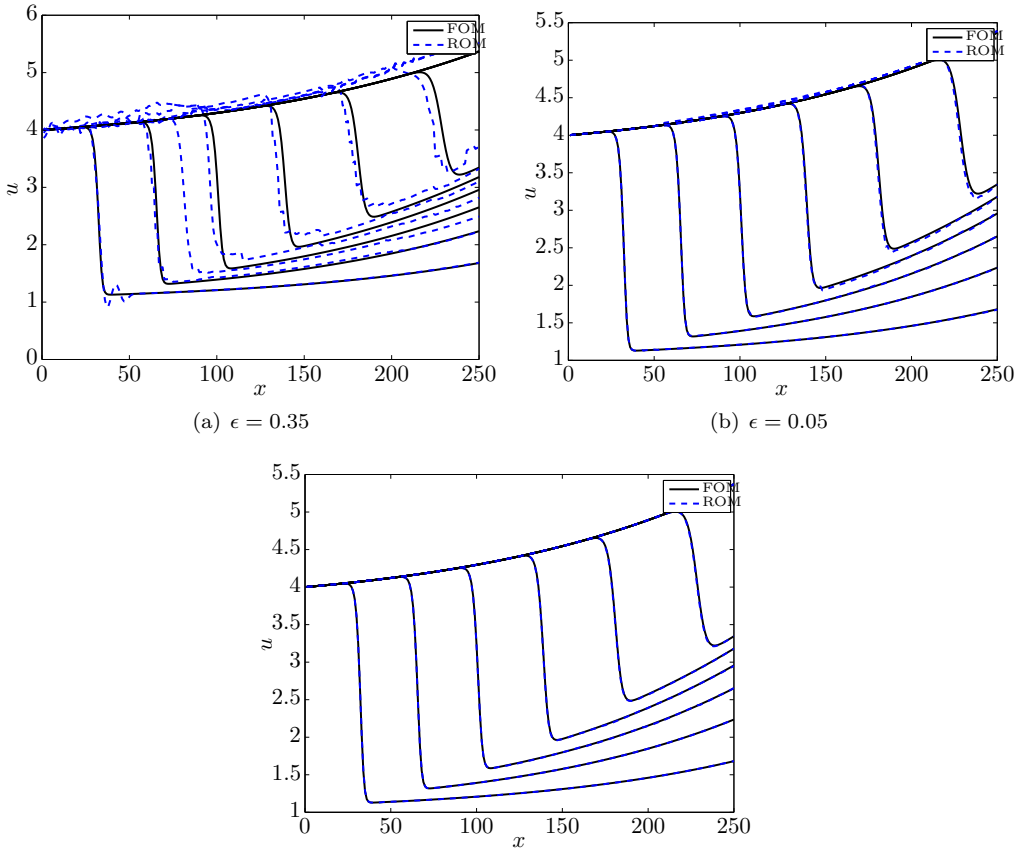

(a) $\epsilon = 0.35$

(b) $\epsilon = 0.05$

(c) $\epsilon = 0.01$

Figure 4: Comparison of solutions computed by h-adaptive POD–Galerkin for different full-order-model tolerances ϵ for the fixed-inputs case.

Greedy [39]) sampling methods, which would lead to a more robust initial basis $\boldsymbol{V}^{(0)}$ and reduce the burden of h-adaptivity to generate accurate results.

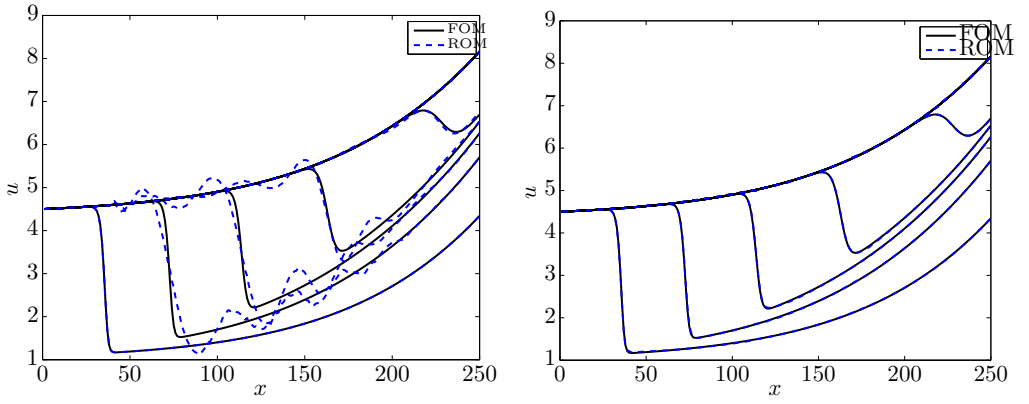

(a) no adaptivity, no truncation ($p = 150$)

(b) h-adaptivity, no truncation ($p^{(0)} = 20$), $n_{\text{reset}} = 100$

Figure 5: Comparison of solutions computed by POD–Galerkin with and without adaptivity for the varying-inputs case.

	no adaptivity			h-adaptivity				
initial basis dimension $p^{(0)}$	10	78	150	5	20	30	20	20
basis-reset frequency n_{reset}				100	100	100	200	50
average basis dimension per Newton iteration $\bar{p}$	10	78	150	69.8	77.2	87.6	130.6	65.6
average number of Refine calls per time step				0.20	0.072	0.07	0.044	0.11
relative error (%)	41.8	1.7	1.4	0.22	0.14	0.45	0.53	0.70
online time (seconds)	1.75	3.54	8.55	6.41	6.06	8.11	9.11	8.78

Table 4: Comparison between POD–Galerkin ROMs without refinement and with h-adaptive refinement for the input-variation case.

7. Conclusions

This work has presented an adaptive h-refinement method for reduced-order models. Key components include 1) an h-refinement mechanism based on basis splitting and tree structure constructed via k-means clustering, 2) dual-weighted residual error indicators, and 3) an adaptive algorithm to moderate when and how to perform the refinement. In contrast to existing *a priori* adaptive methods, the proposed technique provides a mechanism to improve the ROM solution *a posteriori*. As opposed to existing *a posteriori* methods, the proposal does so without incurring any large-scale operations. Numerical examples on the inviscid Burgers equation highlighted the method's ability to accurately predict phenomena not present in the training data used to construct the reduced basis.

Future research directions include incorporating complexity reduction (e.g., empirical interpolation, gappy POD) into the refinement process. In particular, as the reduced basis is refined, sample points (and dual reduced-basis vectors) must be added in a systematic way to ensure the reduced-order model remains solvable. Similar to the manner in which the tree defining the (complete) splitting mechanism is constructed offline, one could generate a hierarchy of these sample points offline from the training data, e.g., by executing [22, Algorithm 3] for n_s equal to the number of nodes in the mesh. In addition, it would be interesting to incorporate a more sophisticated adaptive *coarsening* technique (compared to the simple basis-resetting mechanism in Step 7 of Algorithm 3); for example, one could combine basis vectors whose generalized

coordinates are strongly correlated (or anti-correlated) over recent time steps. Further, it would be interesting to pursue adaptive p-refinement methods, wherein other basis vectors (e.g., truncated POD vectors, discrete wavelets) with possibly global support are added from a library to enrich the reduced basis. In addition, it would be useful to pursue alternative tree-construction methods that satisfy Conditions 1–3 of Section 3.1. Assessing the effect of the proposed refinement method on ROM stability would also constitute an interesting investigation. Finally, it would be advantageous to incorporate Richardson extrapolation in the refinement method to better approximate the outputs of interest; however, this requires knowledge of the convergence rate of the reduced-order model with respect to adding basis vectors.

Appendix A. Refinement algorithm with multiple trees

This section presents a more sophisticated refinement mechanism than that that presented in Section 5. In particular, when a vector is flagged for refinement, it is not necessarily split into all its children. Rather, its children are separated into groups, each of which contributes roughly the same fraction α of the total error for that parent vector. This avoids over-refinement when the number of children is relatively large. However, this leads to an increase in required bookkeeping, as the tree structure changes when children merge: the tree must be altered and separately maintained for each vector. Thus, each basis vector $\boldsymbol{v}_i$, $i = 1, \ldots, p$ will be characterized by its own tree C_i, E_i with m_i nodes, as well as a node on that tree $d_i \in \mathbb{N}(m_i)$.

Algorithm 5 describes the modifications needed to Algorithm 4 to enable this feature. Key modifications include the following. Steps 7–22 separate the children of the parent vector's tree node d_i into groups; the resulting maintenance of the tree structures is performed in Steps 18–19.[5] In steps 23–26, not only is the basis updated, but the trees are as well. Finally, Step 30 performs the necessary bookkeeping for the tree structures due to the removal of redundant basis vectors.

Acknowledgments

The author acknowledges Matthew Zahr for providing the model-reduction testbed that was modified to generate the numerical results, Seshadhri Comandur for helpful discussions related to tree construction and clustering, and the anonymous reviewers for providing insightful remarks and suggestions. This research was supported in part by an appointment to the Sandia National Laboratories Truman Fellowship in National Security Science and Engineering, sponsored by Sandia Corporation (a wholly owned subsidiary of Lockheed Martin Corporation) as Operator of Sandia National Laboratories under its U.S. Department of Energy Contract No. DE-AC04-94AL85000.

References

[1] Benner, P., Gugercin, S., and Willcox, K., "A survey of model reduction methods for parametric systems," *Max Planck Institute Magdeburg Preprints*, Vol. MPIMD/13–14, 2013.

[2] Patera, A. T. and Rozza, G., *Reduced basis approximation and a posteriori error estimation for parametrized partial differential equations*, MIT, 2006.

[3] Gunzburger, M. D., Peterson, J. S., and Shadid, J. N., "Reduced-order modeling of time-dependent PDEs with multiple parameters in the boundary data," *Computer methods in applied mechanics and engineering*, Vol. 196, No. 4, 2007, pp. 1030–1047.

[4] Heinkenschloss, M. and Vicente, L., "Analysis of inexact trust-region SQP algorithms," *SIAM Journal on Optimization*, Vol. 12, No. 2, 2002, pp. 283–302.

[5] Only the lower levels of the tree must be updated, as the current methodology never traverses up a tree.

Algorithm 5 Refine (child grouping) (online)

Input: initial basis $\boldsymbol{V}$, reduced solution $\hat{\boldsymbol{x}}$, child-partition factor $\alpha \leq 1$
Output: refined basis $\boldsymbol{V}$

Compute fine error-estimate vector and fine reduced basis via Algorithm 2:
$\left(\boldsymbol{\delta}^h, \boldsymbol{V}^h\right) \leftarrow$ Error estimates $(\boldsymbol{V}, \hat{\boldsymbol{x}})$.
Put local error estimates in parent–child format $\eta_{ij} = \delta^h_{f(i,j)}$, $i \in \mathbb{N}(p)$, $j \in \mathbb{N}(q_i)$.
Mark basis vectors to refine $I = \{i \mid \sum_j \eta_{ij} \geq 1/p \sum_{kj} \eta_{kj}\}$
for $i \in I$ **do** {Split $\boldsymbol{v}_i$ into k vectors}
  $p \leftarrow \dim(\text{range}(\boldsymbol{V}))$
  Initialize additional-vector count $k \leftarrow 0$ and handled child-node set $D \leftarrow \emptyset$
  while $D \neq \mathbb{N}(q_i)$ **do** {Divide child nodes into groups with roughly equal error}
    $D_k = \arg\min_{z \subset K} \text{card}(z)$, where $K = \{z \subset \mathbb{N}(q_i) \setminus D \mid \sum_{j \in z} \eta_{ij} \geq \alpha \sum_j \eta_{ij}\}$.
    if $D_k = \emptyset$ **then**
      Take all remaining children $D_k = \mathbb{N}(q_i) \setminus D$
    end if
    $\boldsymbol{x}_k = \sum_{j \in D_k} \boldsymbol{v}^h_{f(i,j)}$
    Update tree: $\bar{C}_k \leftarrow C_i$, $\bar{E}_k \leftarrow E_i$
    if $\text{card}(D_k) = 1$ **then** {Use the same tree}
      $\bar{d}_k = C_i(d_i, D_k)$
    else {Alter the tree}
      $\bar{d}_k = d_i$
      $\bar{C}_k\left(\bar{d}_k\right) = \{C_i(d_i, k) \mid k \in D_k\}$
      $\bar{E}_k\left(\bar{d}_k\right) = \bigcup_{k \in \bar{C}_k(d_i)} E_i(k)$
    end if
    $k \leftarrow k + 1$, $D \leftarrow D \cup D_k$
  end while
  $\boldsymbol{v}_i \leftarrow \boldsymbol{x}_0$, $C_i \leftarrow \bar{C}_0$, $E_i \leftarrow \bar{E}_0$ $d_i \leftarrow \bar{d}_0$
  for $l = 1, \ldots, k$ **do**
    $\boldsymbol{v}_{p+l} \leftarrow \boldsymbol{x}_l$, $C_{p+l} \leftarrow \bar{C}_l$, $E_{p+l} \leftarrow \bar{E}_l$, $d_{p+l} \leftarrow \bar{d}_l$
  end for
end for
Compute thin QR factorization with column pivoting $\boldsymbol{V} = \boldsymbol{Q}\boldsymbol{R}$, $\boldsymbol{R}\bar{\boldsymbol{\Pi}} = \bar{\boldsymbol{Q}}\bar{\boldsymbol{R}}$.
Ensure full-rank matrix $\boldsymbol{V} \leftarrow \boldsymbol{V}[\bar{\boldsymbol{\pi}}_1 \ \cdots \ \bar{\boldsymbol{\pi}}_r]$, where r denotes the numerical rank of $\boldsymbol{R}$.
Update tree $[C_1 \ \cdots \ C_r] \leftarrow [C_1 \ \cdots \ C_p][\bar{\boldsymbol{\pi}}_1 \ \cdots \ \bar{\boldsymbol{\pi}}_r]$;
$[E_1 \ \cdots \ E_r] \leftarrow [E_1 \ \cdots \ E_p][\bar{\boldsymbol{\pi}}_1 \ \cdots \ \bar{\boldsymbol{\pi}}_r]$;
$[d_1 \ \cdots \ d_r] \leftarrow [d_1 \ \cdots \ d_p][\bar{\boldsymbol{\pi}}_1 \ \cdots \ \bar{\boldsymbol{\pi}}_r]$.

[5] Eldred, M. S., Weickum, G., and Maute, K., "A multi-point reduced-order modeling approach of transient structural dynamics with application to robust design optimization," *Structural and Multidisciplinary Optimization*, Vol. 38, No. 6, 2009, pp. 599–611.

[6] Arian, E., Fahl, M., and Sachs, E. W., "Trust-region proper orthogonal decomposition for flow control," Tech. Rep. 25, ICASE, 2000.

[7] Ryckelynck, D., "A priori hyperreduction method: an adaptive approach," *Journal of Computational Physics*, Vol. 202, No. 1, 2005, pp. 346–366.

[8] Carlberg, K. and Farhat, C., "An Adaptive POD-Krylov Reduced-Order Model for Structural Optimization," *8th World Congress on Structural and Multidisciplinary Optimization, Lisbon, Portugal*, June 1–5 2009.

[9] Amsallem, D. and Farhat, C., "An Interpolation Method for Adapting Reduced-Order Models and Application to Aeroelasticity," *AIAA Journal*, Vol. 46, No. 7, July 2008, pp. 1803–1813.

[10] Amsallem, D., Cortial, J., Carlberg, K., and Farhat, C., "A method for interpolating on manifolds structural dynamics reduced-order models," *International Journal for Numerical Methods in Engineering*, Vol. 80, No. 9, 2009, pp. 1241–1258.

[11] Eftang, J. L., Patera, A. T., and Rønquist, E. M., "An 'hp' certified reduced basis method for parametrized elliptic partial differential equations," *SIAM Journal on Scientific Computing*, Vol. 32, No. 6, 2010, pp. 3170–3200.

[12] Haasdonk, B., Dihlmann, M., and Ohlberger, M., "A training set and multiple bases generation approach for parameterized model reduction based on adaptive grids in parameter space," *Mathematical and Computer Modelling of Dynamical Systems*, Vol. 17, No. 4, 2011, pp. 423–442.

[13] Drohmann, M., Haasdonk, B., and Ohlberger, M., "Adaptive Reduced Basis Methods for Nonlinear Convection–Diffusion Equations," *Finite Volumes for Complex Applications VI Problems & Perspectives*, Springer, 2011, pp. 369–377.

[14] Peherstorfer, B., Butnaru, D., Willcox, K., and Bungartz, H.-J., "Localized discrete empirical interpolation method," *SIAM Journal on Scientific Computing*, Vol. 36, No. 1, 2014, pp. A168—A192.

[15] Dihlmann, M., Drohmann, M., and Haasdonk, B., "Model reduction of parametrized evolution problems using the reduced basis method with adaptive time partitioning," .

[16] Amsallem, D., Zahr, M. J., and Farhat, C., "Nonlinear model order reduction based on local reduced-order bases," *International Journal for Numerical Methods in Engineering*, Vol. 92, No. 10, December 2012, pp. 891–916.

[17] Barrault, M., Maday, Y., Nguyen, N. C., and Patera, A. T., "An 'empirical interpolation' method: application to efficient reduced-basis discretization of partial differential equations," *Comptes Rendus Mathématique Académie des Sciences*, Vol. 339, No. 9, 2004, pp. 667–672.

[18] LeGresley, P. A., *Application of Proper Orthogonal Decomposition (POD) to Design Decomposition Methods*, Ph.D. thesis, Stanford University, 2006.

[19] Astrid, P., Weiland, S., Willcox, K., and Backx, T., "Missing point estimation in models described by proper orthogonal decomposition," *IEEE Transactions on Automatic Control*, Vol. 53, No. 10, 2008, pp. 2237–2251.

[20] Chaturantabut, S. and Sorensen, D. C., "Nonlinear model reduction via discrete empirical interpolation," *SIAM Journal on Scientific Computing*, Vol. 32, No. 5, 2010, pp. 2737–2764.

[21] Drohmann, M., Haasdonk, B., and Ohlberger, M., "Reduced Basis Approximation for Nonlinear Parametrized Evolution Equations based on Empirical Operator Interpolation," *SIAM Journal on Scientific Computing*, Vol. 34, No. 2, 2012, pp. A937–A969.

[22] Carlberg, K., Farhat, C., Cortial, J., and Amsallem, D., "The GNAT method for nonlinear model reduction: effective implementation and application to computational fluid dynamics and turbulent flows," *Journal of Computational Physics*, Vol. 242, 2013, pp. 623–647.

[23] Barone, M. F., Kalashnikova, I., Segalman, D. J., and Thornquist, H. K., "Stable Galerkin reduced order models for linearized compressible flow," *Journal of Computational Physics*, Vol. 228, No. 6, 2009, pp. 1932–1946.

[24] Lloyd, S., "Least squares quantization in PCM," *Information Theory, IEEE Transactions on*, Vol. 28, No. 2, 1982, pp. 129–137.

[25] Estep, D., "A posteriori error bounds and global error control for approximation of ordinary differential equations," *SIAM Journal on Numerical Analysis*, Vol. 32, No. 1, 1995, pp. 1–48.

[26] Pierce, N. A. and Giles, M. B., "Adjoint recovery of superconvergent functionals from PDE approximations," *SIAM review*, Vol. 42, No. 2, 2000, pp. 247–264.

[27] Babuška, I. and Miller, A., "The post-processing approach in the finite element method—part 1: Calculation of displacements, stresses and other higher derivatives of the displacements," *International Journal for numerical methods in engineering*, Vol. 20, No. 6, 1984, pp. 1085–1109.

[28] Becker, R. and Rannacher, R., *Weighted a posteriori error control in finite element methods*, Vol. preprint no. 96-1, Universitat Heidelberg, 1996.

[29] Rannacher, R., "The dual-weighted-residual method for error control and mesh adaptation in finite element methods," *MAFELEAP*, Vol. 99, 1999, pp. 97–115.

[30] Bangerth, W. and Rannacher, R., *Adaptive finite element methods for differential equations*, Springer, 2003.

[31] Venditti, D. and Darmofal, D., "Adjoint error estimation and grid adaptation for functional outputs: Application to quasi-one-dimensional flow," *Journal of Computational Physics*, Vol. 164, No. 1, 2000, pp. 204–227.

[32] Venditti, D. A. and Darmofal, D. L., "Grid adaptation for functional outputs: application to two-dimensional inviscid flows," *Journal of Computational Physics*, Vol. 176, No. 1, 2002, pp. 40–69.

[33] Park, M. A., "Adjoint-based, three-dimensional error prediction and grid adaptation," *AIAA journal*, Vol. 42, No. 9, 2004, pp. 1854–1862.

[34] Lu, J. C.-C., *An a posteriori error control framework for adaptive precision optimization using discontinuous Galerkin finite element method*, Ph.D. thesis, Massachusetts Institute of Technology, 2005.

[35] Fidkowski, K. J., *A simplex cut-cell adaptive method for high-order discretizations of the compressible Navier-Stokes equations*, Ph.D. thesis, Massachusetts Institute of Technology, 2007.

[36] Meyer, M. and Matthies, H., "Efficient model reduction in non-linear dynamics using the Karhunen-Loève expansion and dual-weighted-residual methods," *Computational Mechanics*, Vol. 31, No. 1, 2003, pp. 179–191.

[37] Drohmann, M. and Carlberg, K., "The ROMES method for reduced-order-model uncertainty quantification," *arXiv preprint 1405.5170*, 2014.

[38] Rewienski, M. J., *A Trajectory Piecewise-Linear Approach to Model Order Reduction of Nonlinear Dynamical Systems*, Ph.D. thesis, Massachusetts Institute of Technology, 2003.

[39] Haasdonk, B. and Ohlberger, M., "Reduced basis method for finite volume approximations of parametrized linear evolution equations," *ESAIM-Mathematical Modelling and Numerical Analysis*, Vol. 42, No. 02, 2008, pp. 277–302.